Submitted to *INFORMS Journal on Computing*
manuscript (Please, provide the mansucript number!)

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

arXiv:1401.0843v1 [math.OC] 4 Jan 2014

Least Squares Policy Iteration with Instrumental Variables vs. Direct Policy Search: Comparison Against Optimal Benchmarks Using Energy Storage

Warren R. Scott
Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544,
wscott@princeton.edu

Warren B. Powell, Somayeh Moazeni
Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544,
powell@princeton.edu, somayeh@princeton.edu

This paper studies approximate policy iteration (API) methods which use least-squares Bellman error minimization for policy evaluation. We address several of its enhancements, namely, Bellman error minimization using instrumental variables, least-squares projected Bellman error minimization, and projected Bellman error minimization using instrumental variables. We prove that for a general discrete-time stochastic control problem, Bellman error minimization using instrumental variables is equivalent to both variants of projected Bellman error minimization. An alternative to these API methods is direct policy search based on knowledge gradient. The practical performance of these three approximate dynamic programming methods are then investigated in the context of an application in energy storage, integrated with an intermittent wind energy supply to fully serve a stochastic time-varying electricity demand. We create a library of test problems using real-world data and apply value iteration to find their optimal policies. These benchmarks are then used to compare the developed policies. Our analysis indicates that API with instrumental variables Bellman error minimization prominently outperforms API with least-squares Bellman error minimization. However, these approaches underperform our direct policy search implementation.

Key words: dynamic programming, approximate dynamic programming, approximate policy iteration, Bellman error minimization, direct policy search, energy storage

History:

1. Introduction

We have long recognized that the powerful theory of discrete Markov decision processes (Puterman 1994) is limited by the well known curse of dimensionality, which can be traced to the need to compute the value of being in each discrete state or, more problematically,

the probability of transitioning from one discrete state to another given an action. Known as a *flat representation* in computer science, even small problems quickly blow up using this model. The field of approximate dynamic programming (Bertsekas and Tsitsiklis (1996b), Powell (2011), Bertsekas (2012)) and reinforcement learning (Sutton and Barto (1998), Szepesvari (2010)) have offered the hope of partially overcoming this problem by replacing the value function using a statistical approximation (in particular, linear regression), allowing us to draw on approximate versions of powerful algorithmic strategies such as value iteration and policy iteration (Puterman 1994).

Somewhat surprisingly, while we have found in our own work that approximate value iteration works well for very specific problem classes (Topaloglu and Powell (2006), Powell et al. (2011), Simao et al. (2009), Nascimento and Powell (2013)), it does not work as a general algorithmic strategy (our success with AVI has always been in the context of problems where we could exploit convexity). The computer science community has focused considerable efforts on a strategy called Q-learning, which involves learning the value of a state-action pair rather than just the value of being in a state (Sutton and Barto (1998)). Q-learning enjoys rigorous convergence theory for lookup table representations (Tsitsiklis 1994), but this does not scale even to small problems. Approximating Q-factors with linear models is more art than science. Convergence results (Sutton et al. (2009b), Maei et al. (2009)) do not contain any performance guarantees (they only establish convergence of a particular approximating architecture), and we are unaware of any empirical comparisons against optimal benchmarks.

Perhaps for this reason, approximate policy iteration has attracted considerable recent attention (see Bertsekas (2012) and the review in Powell and Ma (2011)). API avoids the need to approximate the value of a state-action pair, and enjoys stronger convergence theory, although this always involves assumptions that are unlikely to be perfectly satisfied in practice. However, we are still unaware of any comparisons against optimal benchmarks. In this paper, we use the setting of a class of energy storage problems that requires balancing power from the grid and power from a stochastic, renewable source, to serve a load (that is sometimes time varying), with access to an energy storage device. While the datasets were carefully generated, the focus of the paper is purely algorithmic. Energy storage is a problem that is well suited to the use of value function approximations. Further, while the problem is realistic, we could, with minor simplifications, create

problems that could be solved optimally, providing us with rigorous benchmarks that allow us to provide an accurate assessment of the quality of the solution produced by different algorithmic strategies. The software containing these benchmark problems is available at http://www.castlelab.princeton.edu/datasets.htm.

In this paper, we focus on the use of linear architectures for approximating the value function, a strategy that has received the most attention in the literature. We model our problem in steady state, which allows us to use a powerful algorithmic strategy called Least Squares Approximate Policy Iteration (LSAPI), introduced by Lagoudakis and Parr (2003a), building on the least squares temporal difference method, introduced by Bradtke and Barto (1996) to estimate the value of a fixed policy. Lagoudakis and Parr (2003a) uses a statistical technique that introduces the idea of using instrumental variables (IV) to overcome errors in the observation of the independent variables (explained below).

Recent research has revealed the importance of using the idea of modifying Bellman's optimality equation by first projecting the value function onto the space of linear architectures. This is referred to in the literature as *Mean-Squared Projected Bellman Error* (MSPBE) minimization (see Lagoudakis and Parr (2003b), Sutton et al. (2009a), Maei et al. (2010), and the explanation of this material in Powell (2011)[Chapter 9]). In this paper, we show for the first time that MSPBE is mathematically equivalent to the use of instrumental variables, which is also simpler.

Although our focus is on using exact benchmarks to derive insights into different algorithmic strategies, our choice of energy storage is motivated by the importance of this problem class. For example, Sørensen (1981) addresses the potential benefits of combining a hydro storage facility with wind energy. Three control policies and their corresponding (expected) gained revenues are analyzed in (Barton and Infield 2004), assuming probabilistic models for load and wind, and a load-price curve for the electricity price. Carmona and Ludkovski (2005) and Lai et al. (2010) investigate the operation and value of a gas storage device. Greenblatt et al. (2007) and Swider (2007) discuss incorporating a compressed air energy storage with wind power generation in energy systems. Maximization of the expected market profit of a wind farm and a hydro pumped storage owner over a finite horizon to comply with his commitments in the market is addressed in (Gonzalez et al. 2008). The problem is formulated as a two-stage stochastic programming problem with random market prices and random wind generation. Sioshansi et al. (2009) studies the

potential value of a storage device in the PJM network used for arbitrage. For a thorough review of this growing literature, we refer the reader to (Harsha and Dahleh 2012) and the references therein. The paper that addresses a problem closest to ours is Xi et al. (2014), but this paper is able to use classical backward dynamic programming using lookup tables, and does not explore approximation strategies that might scale to more complex problems.

For our energy storage application, we created a library of test problems using realistic data, constrained by the goal of creating optimal benchmarks. Optimal policies for these problems are computed using exact value iteration (Puterman 1994). These benchmarks are then used to compare variations policies estimated using approximate policy iteration and direct policy search. Our analysis indicates that API with Bellman error minimization using instrumental variables significantly outperforms API with least-squares Bellman error minimization.

This paper makes the following contributions. 1) We introduce and prove the consistency of a hybrid policy based on least squares approximate policy iteration that combines instrumental variables and projected Bellman error minimization. 2) We establish, for the first time, the equivalence of least-squares approximate policy iteration using instrumental variables with mean-squared projected Bellman error minimization, as well as the new hybrid policy combining both instrumental variables and projected Bellman error minimization. 3) We develop a set of benchmark problems using the important context of energy storage where we are able to derive optimal policies. While these problems are relatively simple, CPU times to estimate these policies typically ranged around two *weeks*. 4) We calculated approximate policies using basic LSAPI, LSAPI using instrumental variables, and direct policy search, and compared the results against optimal benchmark policies. These experiments show that LSAPI with instrumental variables performs significantly better than basic LSAPI, and yet even this advanced algorithmic strategy falls far short of optimal. Direct policy search outperforms both LSAPI policies by a wide margin. Since the structure of the policy is the same, the issue is not the accuracy of the approximating architecture, but rather the estimation procedure, calling into question the validity of Bellman error minimization.

This paper is organized as follows. Section 2 provides an overview of approximate policy iteration methods. Several policy evaluation algorithms based on Bellman error minimization are discussed in Section 3. The proof for their equivalence is presented in Section 4.

Direct policy search and knowledge gradient are discussed in Section 5. The energy storage management problem, its underlying stochastic processes, and its stochastic dynamic programming formulation are explained in Section 6. Performance of the computed approximate dynamic programming policies is investigated and compared with benchmark problems in Section 7. Concluding remarks are given in Section 8.

2. Approximate Policy Iteration Algorithms

Stochastic dynamic programming for maximizing expected revenues relies on Bellman's optimality equation given by

$$V(S_t) = \max_x \ \mathbb{E}\left[C\left(S_t, x\right) + \gamma \ V\left(S_{t+1}\right) | S_t\right], \tag{1}$$

where S_t refers to the state variable at time step t, $V(\cdot)$ is the optimal value function (around the pre-decision state S_t), $C(S_t, x)$ is the contribution function, and $0 \leq \gamma < 1$ is the discount factor. The expectation in equation (1) is over the exogenous random changes, denoted by W_{t+1}, in the state of the system. The state variable at the next step is then obtained by the transition function $S_{t+1} = S^M\left(S_t, x, W_{t+1}\right)$. Throughout, we use the convention that any variable indexed by t is known at time t.

Often computing the expected value of $V(S_t)$ is challenging and needs to be approximated. To avoid this, a modified version of Bellman's equation based on the *post-decision state variables* can be adopted, see e.g., (Bertsekas 2011) or Section 4.6 of (Powell 2011) for a thorough discussion of post-decision states. The post-decision state variable, denoted by S_t^x, refers to the state immediately after being in the pre-decision state S_t and taking the action/decision x, but before any exogenous information (randomness) from the state transition has been revealed. The *post-decision value function*, denoted by $V^x(S_t^x)$, explains the value of being in the post-decision state S_t^x. The relationship between $V^x(S_t^x)$ and $V(S_t)$ is defined as:

$$V^x\left(S_t^x\right) \stackrel{\text{def}}{=} \mathbb{E}\left[V\left(S_{t+1}\right) \ | \ S_t^x\right].$$

Thus, Bellman's equation (1) around the post-decision state variables can be written as

$$V_{t-1}^x(S_{t-1}^x) = \mathbb{E}\left[\max_x \ \left(C\left(S_t, x\right) + \gamma \ V_t^x\left(S_t^x\right)\right) \ | \ S_{t-1}^x\right]. \tag{2}$$

The expectation being outside of the maximum operator allows us to solve the inner maximization problem using deterministic optimization techniques. Furthermore, in some problems, the post-decision state has a lower dimension than the pre-decision state.

For most applications, however, when the state variable is multidimensional and continuous, Bellman's equations (1) or (2) cannot be solved exactly; as a result a large field of research on approximation techniques has evolved, see e.g., (Bertsekas and Tsitsiklis 1996a, Sutton and Barto 1998, Szepesvari 2010, Powell 2011). In this paper, we focus on the widely used strategy of linear architectures of the form

$$\overline{V}_t^x(S_t^x) \stackrel{\text{def}}{=} \sum_{f\in\mathcal{F}} \theta_f \ \phi_f(S_t^x) = \theta^\top \phi(S_t^x), \tag{3}$$

where $\phi(S_t^x)$ is the column vector with elements $\{\phi_f(S_t^x)\}_{f\in\mathcal{F}}$, and θ is the column vector of weights for the basis functions. Substituting approximate post-decision value functions (3) into equation (2), for a fixed policy π and a given vector θ, we get

$$\theta^\top \phi\left(S_{t-1}^x\right) = \mathbb{E}\left[C\left(S_t, X^\pi(S_t|\theta)\right) + \gamma\ \theta^\top \phi(S_t^x) \ \mid\ S_{t-1}^x\right]. \tag{4}$$

In general, a fixed point satisfying this equation does not exist using linear architectures (de Farias and Van Roy 2000), but this equation forms the foundation of LSAPI. For a given weight vector θ, our policy $X^\pi(S_t|\theta)$ is given by

$$X^\pi(S_t|\theta) = \text{argmax}_x \ \left[C(S_t, x) + \gamma\ \theta^\top \phi(S_t^x)\right]. \tag{5}$$

The basic LSAPI algorithm of Bradtke and Barto (1996) is given in figure 1. We use the adaption for the post-decision state given in Ma and Powell (2010). Let $\{S_{t-1,i}^x\}_{i=1}^N$ be a set of (randomly generated) post-decision states. For each sample $i = 1, \cdots, N$, the next post-decision state $S_{t,i}^x$ is simulated and, using equation (5), the the optimal decision $X^\pi(S_{t,i}|\theta)$ is computed. For every $i = 1, \cdots, N$, equation (4) can be rewritten as:

$$\begin{aligned} C\left(S_{t,i}, X^\pi(S_{t,i}|\theta)\right) = &\left(\phi(S_{t-1,i}^x) - \gamma\mathbb{E}[\phi(S_{t,i}^x)|S_{t-1,i}^x]\right)^\top \theta \\ &+ C(S_{t,i}, X^\pi(S_{t,i}|\theta)) - \mathbb{E}[C(S_{t,i}, X^\pi(S_{t,i}|\theta))|S_{t-1,i}^x]. \end{aligned} \tag{6}$$

Denote $C_{t,i} \stackrel{\text{def}}{=} C\left(S_{t,i}, X^\pi(S_{t,i}|\theta)\right)$ and $X_{t,i} \stackrel{\text{def}}{=} \left(\phi(S_{t-1,i}^x) - \gamma\mathbb{E}[\phi(S_{t,i}^x)|S_{t-1,i}^x]\right)^\top$. Using simulation, we are able to get observations of $C_{t,i}$ and $X_{t,i}$ in equation (6).

Approximate Policy Iteration Algorithm
(01) Initialize θ.
(02) **for** $j=1$ **to** M **(Policy Improvement Loop)**
(03) **for** $i=1$ **to** N **(Policy Evaluation Loop)**
(04) Simulate a random post-decision state, $S^x_{t-1,i}$.
(05) Record $\phi(S^x_{t-1,i})$.
(06) Simulate the state transition to get $S_{t,i}$.
(07) Determine the decision, $x=X^\pi(S_{t,i}\|\theta)$ using equation (5).
(08) Record $C(S_{t,i},x)$.
(09) Record $\phi(S^x_{t,i})$, the observation of $\mathbb{E}[\phi(S^x_{t,i})\|S^x_{t-1,i}]$.
(10) **End**
(11) Update θ using equation (9), or (10), or (12), or (13). **(Policy Evaluation)**
(12) **End**

Figure 1 **Summary of (steady-state) approximate policy iteration. Policy Evaluation loop approximately evaluates the cost function of a given policy. Policy Improvement loop generates a new (improved) policy. The constants M and N indicate the maximum number of policy improvement and policy evaluation iterations.**

Let Φ_t and Φ_{t-1} be the matrices of fixed basis functions whose rows correspond to states and columns correspond to basis functions. Furthermore, refer to the vector of contribution functions and its expected value by C_t and $\overline{C}_t$, respectively. More precisely,

$$\Phi_{t-1} \stackrel{\text{def}}{=} \begin{bmatrix} \phi(S^x_{t-1,1})^\top \\ \vdots \\ \phi(S^x_{t-1,N})^\top \end{bmatrix}, \quad \Phi_t \stackrel{\text{def}}{=} \begin{bmatrix} \phi(S^x_{t,1})^\top \\ \vdots \\ \phi(S^x_{t,N})^\top \end{bmatrix}, \quad C_t \stackrel{\text{def}}{=} \begin{bmatrix} C_{t,i} \\ \vdots \\ C_{t,N} \end{bmatrix}, \quad \overline{C}_t \stackrel{\text{def}}{=} \begin{bmatrix} \overline{C}_{t,1} \\ \vdots \\ \overline{C}_{t,N} \end{bmatrix},$$

where $\overline{C}_{t,i} \stackrel{\text{def}}{=} \mathbb{E}[C(S_{t,i},X^\pi(S_{t,i}|\theta))|S^x_{t-1,i}]$. Thus, equation (6) can be presented in a matrix form as

$$\underbrace{C_t}_{n\times 1} = \underbrace{(\Phi_{t-1}-\gamma\Phi_t)}_{n\times k}\underbrace{\theta}_{k\times 1} + \underbrace{C_t-\overline{C}_t}_{n\times 1}, \tag{7}$$

We refer to $(C_t-\overline{C}_t)$ as the *Bellman errors*, although this is not the standard definition of the Bellman error.

An improved value for θ can be obtained by minimizing the Euclidean norm of the Bellman errors,

$$\min_{\theta} \ \left\|C_t - \overline{C}_t\right\|_2^2, \tag{8}$$

which is the goal in least-squares Bellman error minimization. In the next section, we address a few of its variants including Bellman error minimization using instrumental variables and projected Bellman error minimization, after which we prove that these two methods are equivalent. In our discussion, we make the following assumption about the linear independence of the basis functions and underlying matrices.

Assumption 2.1 *The matrices* Φ_{t-1}, $(\Phi_{t-1}-\gamma\Phi_t)$, *and* $\Phi_{t-1}^{\top}(\Phi_{t-1}-\gamma\Phi_t)$ *have full column rank, and* $k \leq n$.

This assumption ensures that enough different states are visited that allows us to identify the linear model.

3. Least Squares Temporal Differences (LSTD) Policy Evaluation

The LSTD method, introduced by Bradtke and Barto (1996), is one of the (batch) variants of the TD algorithms. For a given policy and assuming a linear model with fixed basis functions for the value function, this on-policy algorithm performs a least squares regression so that the sum of the temporal differences over the simulation to approximate the expectation equals zero. LSTD tends to be statistically efficient and extract more information from training experiences, compared to some other TD methods. The method of instrumental variables, see e.g., Young (2011), is used in LSTD to compensate for the use of Monte Carlo simulation to approximate the expected cost in $\overline{C}_{t,i} \stackrel{\text{def}}{=} \mathbb{E}[C(S_{t,i}, X^{\pi}(S_{t,i}|\theta))|S^x_{t-1,i}]$. The estimation of the value function converges with probability one to the true value, when the true value function lies in the span of the basis functions, the number of state transitions increases to infinity, and a few other technical assumptions are fulfilled.

Lagoudakis and Parr (2003b) expands upon the LSTD algorithm in (Bradtke and Barto 1996) by using a linear architecture to approximate the value function over the higher dimension state-action pairs. Furthermore, they give a geometric interpretation of several different methods of approximately solving Bellman's equation. Also assuming finite states and actions, Lagoudakis and Parr (2003b) introduces least-squares policy iteration which

approximates the value of state-action pairs (Q-factors) with a linear model. When doing policy evaluation, they choose to use least-squares to minimize the fixed-point approximation error instead of the Bellman error. An alternative approach is to project the Bellman errors down onto the space spanned by the basis functions of the value function and then minimize the Bellman errors. Lagoudakis and Parr (2003b) explains that in general the approximate value function is a fixed point of the projected Bellman operator, not the Bellman operator (see De Farias and Van Roy (2000) for a nice discussion).

Bradtke and Barto (1996) addresses conditions for on-policy policy evaluation subproblem based on Bellman error minimization to converge to a fixed value function when using a linear model for the value function. Given a fixed policy, Tsitsiklis and Van Roy (1997) provide conditions under which on-policy temporal-difference learning will converge to the true value function projected onto the space of value function approximations (with respect to a weighted norm). Furthermore, Tsitsiklis and Van Roy (1997) proves that, when evaluating a fixed policy, off-policy TD(λ) algorithms with a linear function approximation of the value function may not converge. In this case, off-policy means that the distribution of the states visited during a single infinite trajectory is not equal to the distribution of the states visited if we followed the fixed policy of interest. However, Lagoudakis and Parr (2003b) explain that on-policy LSTD biases the value function and may fail to acceptably fit the value function at states that are rarely visited. Another major disadvantage of on-policy learning is that Assumption 2.1 may not hold if the policy does not explore enough states. An important point to keep in mind is that the value of the greedy policy determined by the final value function approximation may be significantly different from the true value function.

In the following, we summarize the policy evaluation algorithms based on Bellman error minimization and instrumental variables. We focus on the off-policy case in both subsections (3.1) and (3.2). Subsection 3.1 addresses least-squares Bellman error minimization and instrumental variables Bellman error minimization. Subsection 3.2 explains least-squares projected Bellman error minimization and instrumental variables projected Bellman error minimization.

3.1. Policy Evaluation using Bellman Error Minimization

Applying the typical method of least-squares, a solution of equation (8) equals,

$$\hat{\theta} = \left[(\Phi_{t-1} - \gamma\Phi_t)^\top(\Phi_{t-1} - \gamma\Phi_t)\right]^{-1}(\Phi_{t-1} - \gamma\Phi_t)^\top C_t, \tag{9}$$

which we refer to as the *least-squares Bellman error minimization* estimator. The matrix of regressors, $(\Phi_{t-1} - \gamma\Phi_t)$, is not deterministic ($\Phi_t$ is not deterministic because we cannot calculate $\mathbb{E}[\phi(S_t^x)|S_{t-1}^x]$); we can only simulate $\phi(S_t^x)$ given S_{t-1}^x and, as a result, the least-squares estimator for θ will typically be inconsistent.

A class of simple and computationally efficient techniques to obtain consistent estimates, without modeling the noise, is referred to as *Instrumental Variable Methods* (IVM). An instrumental variable is a variable that is correlated with the regressors, but uncorrelated with the errors in the regressors and the observations, see e.g., (Durbin 1954, Kendall and Stuart 1961, Söderström and Stoica 1983, Bowden and Turkington 1984). We provide a brief overview of IV methods in Appendix A. Instrumental variables have been previously used in the context of approximate policy iteration algorithms, see e.g., (Bradtke and Barto 1996, Ma and Powell 2010), but otherwise have not received much attention, even in the reinforcement learning literature.

This results in what we call *instrumental variables Bellman error minimization*,

$$\hat{\theta} = \left[(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)\right]^{-1}(\Phi_{t-1})^\top C_t. \tag{10}$$

Bradtke and Barto (1996) provide conditions such that equation (10) is a consistent estimator ($\lim_{n\to\infty}\hat{\theta} = \theta$ with probability one) for the on-policy case. The proof references the consistency properties of the method of instrumental variables by showing that the columns of Φ^n are appropriate instrumental variables (see Appendix A). One interesting comment is that the matrix $[\Phi_{t-1}^\top(\Phi_{t-1} - \gamma\Phi_t)]$ could have negative eigenvalues, unlike $\left[(\Phi_{t-1} - \gamma\Phi_t)^\top(\Phi_{t-1} - \gamma\Phi_t)\right]$.

3.2. Policy Evaluation using Projected Bellman Error Minimization

The idea of projected Bellman error minimization, also called *least-squares fixed-point approximation* in Lagoudakis and Parr (2003b), is to first project the Bellman errors onto the space spanned by the basis functions defining the value function and then minimize them, see (Lagoudakis and Parr 2003b, Sutton et al. 2009c). Projecting the left and right hand sides of equation (7) onto the space spanned by Φ_{t-1} with respect to the Euclidean norm, we get

$$\Pi_{t-1}C_t = \Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\theta + \Pi_{t-1}(C_t - \overline{C}_t), \tag{11}$$

where $\Pi_{t-1} = \Phi_{t-1}((\Phi_{t-1})^\top \Phi_{t-1})^{-1}(\Phi_{t-1})^\top$ is the projection operator onto the space spanned by the basis functions (see (Tsitsiklis and Van Roy 1997) for the original derivation of this mapping or Section 8.2.3 of Powell (2011)). Assumption 2.1 implies that the matrix Φ_{t-1} has full column rank, and hence Π_{t-1} is well defined. We refer to $\Pi_{t-1}(C_t - \overline{C}_t)$, as the *projected Bellman error*. Notice that the $n \times n$ matrix Π_{t-1} is the projection with respect to the Euclidean norm that solves the problem: $\min_\Pi \|\Phi_{t-1}\theta - b\|_2 = \|\Pi_{t-1}b - b\|_2$. In other words, to get from an arbitrary vector b to the closest vector (in the Euclidean norm sense) which is in the span of the columns of Φ_{t-1}, it is sufficient to just apply the projection Π_{t-1} to the vector b.

Typically, equation (11) is an over-determined set of equations. Taking a least squares approach, we find θ by minimizing the norm of the projected Bellman error

$$\min_\theta \ \left\|\Pi_{t-1}(C_t - \overline{C}_t)\right\|_2 = \min_\theta \ \left\|\Pi_{t-1}C_t - \Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\theta\right\|_2 .$$

The least-squares estimator of θ then yields

$$\hat{\theta} = \left[\left(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\right)^\top \left(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\right)\right]^{-1} \left(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\right)^\top \Pi_{t-1}C_t, \quad (12)$$

to which we refer as *least-squares projected Bellman error minimization*.

To establish a consistent estimator for θ, similar to the derivation of equation (10), Φ_{t-1} can be used as an instrumental variable, e.g., see Appendix A or the proof in (Bradtke and Barto 1996). We refer to the resulting estimator as the *instrumental variables projected Bellman error minimization*,

$$\hat{\theta} = \left[(\Phi_{t-1})^\top \Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\right]^{-1} (\Phi_{t-1})^\top \Pi_{t-1}C_t. \quad (13)$$

We note that $\Pi_{t-1}\Phi_{t-1}$ could also have been used for the instrumental variables instead of Φ_{t-1}. However, it is easy to see that the obtained estimator would be equivalent to that in equation (13).

In the rest of this section, we show that the estimator $\hat{\theta}$ in equation (13) is consistent (converges in probability to the true weights), for the off-policy case. Bradtke and Barto (1996) proves a similar result for the on-policy case. The following discussion remains valid even when the state space is continuous or the discount factor is one.

We first make the following assumptions:

Assumption 3.1 *The rows of Φ_{t-1} are i.i.d. (off-policy).*

Assumption 3.2 *The covariance matrix Σ between the instrumental variables and regressors has full rank k, where $\Sigma_{jl} = Cov[(\Phi_{t-1})_j,\ X_l]$ and $X = \Pi_{t-1}(\Phi_{t-1} - \gamma\mathbb{E}[\gamma\Phi_t|\{S^x_{t-1}\}])$.*

Assumption 3.3 $\mathbb{E}[|(\Phi_{t-1})_{ij}Y^{''}_i|] < \infty$, *for $j = 1, 2, \cdots, k$, where $Y'' = \Pi_{t-1}(C_t - \overline{C}_t)$.*

Assumption 3.4 $\mathbb{E}[|(\Phi_{t-1})_{ij}\left(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t) - X\right)_{il}|] < \infty$, *for $j, l \in \{1, 2, \cdots, k\}$, where the matrix X as in Assumption* (3.2).

Assumptions 3.3 and 3.4 are necessary to use the law of large numbers, which guarantees sample means converge to their true means.

Theorem 3.1 *Let assumptions 2.1, 3.1, 3.2, 3.3, and 3.4 hold. The estimator $\hat{\theta}$ presented in equation* (13) *is a consistent estimator for θ defined in equation* (11).

Proof: *The proof follows from Theorem A.1 in Appendix A. The following lemmas show that the assumptions in Appendix A for Theorem A.1 hold. For notation consistent with Appendix A, we let $X \overset{\text{def}}{=} \Pi_{t-1}(\Phi_{t-1} - \gamma\mathbb{E}[\gamma\Phi_t|\{S^x_{t-1}\}])$, $X' \overset{\text{def}}{=} \Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)$, $X'' \overset{\text{def}}{=} X' - X$, $Y'' \overset{\text{def}}{=} \Pi_{t-1}(C_t - \overline{C}_t)$, and $Z \overset{\text{def}}{=} \Phi_{t-1}$. We first show Assumption A.1 holds.*

Lemma 3.1 *For every i, $\mathbb{E}[Y^{''}_i] = 0$.*

Proof:

$$\begin{aligned}\mathbb{E}[Y^{''}] &= \mathbb{E}[\Pi_{t-1}(C_t - \overline{C}_t)] \\ &= \mathbb{E}\left[\mathbb{E}[\Pi_{t-1}(C_t - \overline{C}_t)|\{S^x_{t-1}\}]\right] \\ &= \mathbb{E}[\Pi_{t-1}\underbrace{\mathbb{E}[(C_t - \overline{C}_t)|\{S^x_{t-1}\}]}_{=\vec{0}}] = \vec{0}.\end{aligned}$$

We next show Assumption A.2 holds, which states that the mean of the noise in the observation of the explanatory variables is zero.

Lemma 3.2 *For every i, j, $\mathbb{E}[X^{''}_{ij}] = 0$.*

Proof:

$$\begin{aligned}\mathbb{E}[X^{''}] &= \mathbb{E}[X^{'} - X] \\ &= \mathbb{E}\left[\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t) - \Pi_{t-1}(\Phi_{t-1} - \mathbb{E}[\gamma\Phi_t|\{S^x_{t-1}\}])\right]\end{aligned}$$

$$
\begin{aligned}
&= \gamma\mathbb{E}\left[\Pi_{t-1}(\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \Phi_t)\right] \\
&= \gamma\mathbb{E}\left[\mathbb{E}[\Pi_{t-1}(\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \Phi_t)|\{S^x_{t-1}\}]\right] \\
&= \gamma\mathbb{E}\left[\Pi_{t-1}\mathbb{E}[\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \Phi_t|\{S^x_{t-1}\}]\right] \\
&= \gamma\mathbb{E}[\Pi_{t-1}(\underbrace{\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \mathbb{E}[\Phi_t|\{S^x_{t-1}\}]}_{=\mathbf{0}})] = 0.
\end{aligned}
$$

We next show Assumption A.3 holds, which states that the instrumental variables are uncorrelated with the noise in the observations of the response variable. Below, e_i denotes the column vector of all zeros except at the ith element which equals 1.

Lemma 3.3 *For every i, j, $Cov[Z_{ij}, Y_i^{''}] = 0$.*

Proof:

$$
\begin{aligned}
Cov[Z_{ij}, Y_i^{''}] &= \mathbb{E}[Z_{ij}Y_i^{''}] - \mathbb{E}[Z_{ij}]\underbrace{\mathbb{E}[Y_i^{''}]}_{=0} \\
&= \mathbb{E}[\Phi_{t-1,ij}\{\Pi_{t-1}(C_t - \overline{C}_t)\}_i] \\
&= \mathbb{E}[\Phi_{t-1,ij}e_i^\top\ \Pi_{t-1}(C_t - \overline{C}_t)] \\
&= \mathbb{E}\left[\mathbb{E}[\Phi_{t-1,ij}e_i^\top\ \Pi_{t-1}(C_t - \overline{C}_t)|\{S^x_{t-1}\}]\right] \\
&= \mathbb{E}[\Phi_{t-1,ij}e_i^\top\ \Pi_{t-1}\underbrace{\mathbb{E}[C_t - \overline{C}_t|\{S^x_{t-1}\}]}_{=\vec{0}}] = 0.
\end{aligned}
$$

We next show Assumption A.4 holds.

Lemma 3.4 *For every i, j, and l, $Cov[Z_{ij}, X_{il}^{''}] = 0$.*

Proof:

$$
\begin{aligned}
Cov[Z_{ij}, X_{il}^{''}] &= \mathbb{E}[Z_{ij}X_{il}^{''}] - \mathbb{E}[Z_{ij}]\underbrace{\mathbb{E}[X_{il}^{''}]}_{=0} \\
&= \mathbb{E}[Z_{ij}(X_{il}^{'} - X_{il})] \\
&= \mathbb{E}[Z_{ij}e_i^\top(X^{'} - X)e_l] \\
&= \gamma\mathbb{E}[\Phi_{t-1,ij}e_i^\top\Pi_{t-1}(\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \Phi_t)e_l] \\
&= \gamma\mathbb{E}[\mathbb{E}[\Phi_{t-1,ij}e_i^\top\Pi_{t-1}(\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \Phi_t)e_l|\{S^x_{t-1}\}]] \\
&= \gamma\mathbb{E}[\Phi_{t-1,ij}e_i^\top\Pi_{t-1}\mathbb{E}[\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \Phi_t|\{S^x_{t-1}\}]e_l] \\
&= \gamma\mathbb{E}[\Phi_{t-1,ij}e_i^\top\Pi_{t-1}\underbrace{(\mathbb{E}[\Phi_t|\{S^x_{t-1}\}] - \mathbb{E}[\Phi_t|\{S^x_{t-1}\}])}_{=0}e_l] = 0.
\end{aligned}
$$

Finally, Assumption A.5 holds by Assumption 3.2, and assumptions A.6, A.7, and A.8 follow trivially from assumptions 3.2, 3.1, 3.3, and 3.4 by the law of large numbers (see Appendix A). Therefore Theorem A.1 follows. Q.E.D.

We proceed by showing that instrumental variables Bellman error minimization, least-squares projected Bellman error minimization, and instrumental variables projected Bellman error minimization methods, are equivalent.

4. Equivalence of Instrumental Variables Bellman Error Minimization and Projected Bellman Error Minimization Methods

The following theorem points out the relationship among the three estimators in equations (10), (12), and (13).

Theorem 4.1 *Under Assumption 2.1, the following policy evaluation algorithms are equivalent:*

- *Instrumental Variables Bellman Error Minimization (LSTD) in equation* (10)*:*

$$\hat{\theta} = [(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top C_t.$$

- *Least-Squares Projected Bellman Error Minimization (Least-Squares Fixed Point Approximation) in equation* (12)*:*

$$\hat{\theta} = \left[(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))^\top (\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))\right]^{-1} (\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))^\top \Pi_{t-1} C_t.$$

- *Instrumental Variables Projected Bellman Error Minimization in equation* (13)*:*

$$\hat{\theta} = \left[(\Phi_{t-1})^\top \Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\right]^{-1} (\Phi_{t-1})^\top \Pi_{t-1} C_t.$$

Proof: *We first show equations* (10) *and* (13) *are equivalent. Starting with equation* (13) *and recalling* $\Pi_{t-1} = \Phi_{t-1}((\Phi_{t-1})^\top\Phi_{t-1})^{-1}(\Phi_{t-1})^\top$*, we can write*

$$\begin{aligned}
&\left[(\Phi_{t-1})^\top \Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\right]^{-1} (\Phi_{t-1})^\top \Pi_{t-1} C_t \\
&= \left[(\Phi_{t-1})^\top (\underbrace{\Pi_{t-1}\Phi_{t-1}}_{\Phi_{t-1}} - \gamma\Pi_{t-1}\Phi_t)\right]^{-1} \underbrace{(\Phi_{t-1})^\top\Phi_{t-1}((\Phi_{t-1})^\top\Phi_{t-1})^{-1}}_{I_{k\times k}} (\Phi_{t-1})^\top C_t \\
&= \left[(\Phi_{t-1})^\top\Phi_{t-1} - \gamma(\Phi_{t-1})^\top\Pi_{t-1}\Phi_t\right]^{-1} (\Phi_{t-1})^\top C_t
\end{aligned}$$

$$
\begin{aligned}
&= \left[(\Phi_{t-1})^\top\Phi_{t-1} - \gamma\underbrace{(\Phi_{t-1})^\top\Phi_{t-1}((\Phi_{t-1})^\top\Phi_{t-1})^{-1}}_{I_{k\times k}}(\Phi_{t-1})^\top\Phi_t\right]^{-1}(\Phi_{t-1})^\top C_t\\
&= \left[(\Phi_{t-1})^\top\Phi_{t-1} - \gamma(\Phi_{t-1})^\top\Phi_t\right]^{-1}(\Phi_{t-1})^\top C_t\\
&= \left[(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)\right]^{-1}(\Phi_{t-1})^\top C_t.
\end{aligned}
$$

Hence equations (10) *and* (13) *are equivalent. Next, we show equations* (10) *and* (12) *are equivalent. We start by writing*

$$
\begin{aligned}
&(\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\\
&= (\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)(\Phi_{t-1} - \gamma\Phi_t)^\top\Phi_{t-1}[(\Phi_{t-1})^\top\Phi_{t-1}]^{-1}(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)\\
&= (\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)\\
\Longrightarrow\quad &(\Phi_{t-1} - \gamma\Phi_t)^\top\Phi_{t-1}[(\Phi_{t-1})^\top\Phi_{t-1}]^{-1}\underbrace{(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)[(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}}_{I_k}(\Phi_{t-1})^\top\\
&= (\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)[(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top\\
\Longrightarrow\quad &(\Phi_{t-1} - \gamma\Phi_t)^\top\underbrace{\Phi_{t-1}[(\Phi_{t-1})^\top\Phi_{t-1}]^{-1}(\Phi_{t-1})^\top}_{\Pi_{t-1}}\\
&= (\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)[(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))^\top\\
&= (\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)[(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top \qquad (14)\\
\Longrightarrow\quad &[(\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))^\top\\
&= \underbrace{[(\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1} - \gamma\Phi_t)^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)}_{I_k}\\
\Longrightarrow\quad &[(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top\\
\Longrightarrow\quad &[(\Phi_{t-1} - \gamma\Phi_t)^\top\underbrace{\Pi_{t-1}}_{(\Pi_{t-1})^\top\Pi_{t-1}}(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))^\top\\
&= [(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top\\
\Longrightarrow\quad &[(\Phi_{t-1} - \gamma\Phi_t)^\top(\Pi_{t-1})^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1} - \gamma\Phi_t)^\top\underbrace{(\Pi_{t-1})^\top}_{(\Pi_{t-1})^\top\Pi_{t-1}}\\
&= [(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top\\
\Longrightarrow\quad &[(\Phi_{t-1} - \gamma\Phi_t)^\top(\Pi_{t-1})^\top\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1} - \gamma\Phi_t)^\top(\Pi_{t-1})^\top\Pi_{t-1}\\
&= [(\Phi_{t-1})^\top(\Phi_{t-1} - \gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top\\
\Longrightarrow\quad &\left[(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))^\top(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))\right]^{-1}(\Pi_{t-1}(\Phi_{t-1} - \gamma\Phi_t))^\top\Pi_{t-1}
\end{aligned}
$$

$$
\begin{aligned}
&= [(\Phi_{t-1})^\top(\Phi_{t-1}-\gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top \\
\implies \quad &\left[(\Pi_{t-1}(\Phi_{t-1}-\gamma\Phi_t))^\top(\Pi_{t-1}(\Phi_{t-1}-\gamma\Phi_t))\right]^{-1}(\Pi_{t-1}(\Phi_{t-1}-\gamma\Phi_t))^\top\Pi_{t-1}C_t \\
&= [(\Phi_{t-1})^\top(\Phi_{t-1}-\gamma\Phi_t)]^{-1}(\Phi_{t-1})^\top C_t.
\end{aligned}
$$

Equation (14) *uses the fact that* $(\Pi_{t-1})^\top = \Pi_{t-1}$. *Hence equations* (10), (12), *and* (13) *are equivalent.* *Q.E.D.*

The results in Section 7 illustrate that the Least-squares Bellman error minimization cannot be equivalent to the Bellman error minimization with instrumental variables, and thus the two others.

5. Direct Policy Search

An alternative to Bellman error minimization for finding the regression vector θ is direct policy search. Consider policies parameterized by θ of the form in equation (5), in which the post-decision value function has been replaced by the linear model $\phi(S^x)^\top\theta$. In contrast to policy iteration or value iteration methods, the goal in direct policy search is not necessarily to find a value function which is close (with respect to some norm) to the true value function; the objective is to find a parameter vector θ for which the policy $X^\pi(s|\theta)$ performs well, i.e., it solves the following stochastic optimization problem

$$\max_\theta F(\theta) = \mathbb{E}\sum_{t=0}^{T}\gamma^t C(S_t, X^\pi(S_t|\theta)), \tag{15}$$

given the policy $X^\pi(S_t|\theta)$. Solving this problem becomes challenging, particularly as the dimension of θ grows. Furthermore, the optimization problem given by equation (15) is typically non-convex and non-separable. Classic stochastic optimization algorithms can be used to sequentially choose policies to simulate. When the dimension of θ is small, the Knowledge Gradient for Continuous Parameters (KGCP) policy has been shown to work well for efficiently optimizing θ, e.g., see (Scott et al. 2011). This approach is explained in the following subsection.

Note that, in direct policy search, we only need to consider features which are a function of the decisions. In addition, $F(\theta)$ is nonconvex, lacks easily computable derivatives and can only be simulated (introducing noise), making the search for an optimal θ fairly difficult for higher dimensional problems. In our experiments, θ is limited to two or three dimensions.

5.1. The Knowledge Gradient for Direct Policy Search

For a given value of θ^i, we can obtain a noisy observation of the objective in equation (15) by simulating

$$\hat{V}^\pi(S_0) = C_0(S_0, X^\pi(S_0|\theta^i)) + \gamma C_1(S_1, X^\pi(S_1|\theta^i)) + \gamma^2 C_2(S_2, X^\pi(S_2|\theta^i)) + \cdots.$$

The KGCP policy for optimizing θ treats the objective function $\mu(\theta)$ as a Gaussian process regression. It then combines a model of $\mu(\theta) \stackrel{\text{def}}{=} \hat{V}^\pi(S)$ with a criterion which chooses the next value of θ for which a noisy observation of $\mu(\theta)$ will be simulated. The KGCP quantifies how much the maximum of the objective is expected to increase by getting an additional noisy observation of $\mu(\theta)$ at a particular value of θ.

More formally, let $\mathcal{F}^n$ be the sigma-algebra generated by $\theta^0, \cdots, \theta^{n-1}$ and the corresponding noisy observations of $\mu(\theta^0), \cdots, \mu(\theta^{n-1})$. Denote the updated expected values of μ at θ^i, conditioned on $\mathcal{F}^n$, by $\mu^n(\theta^i)$. The KGCP is defined as

$$\bar{\nu}^{KG,n}(\theta) \stackrel{\text{def}}{=} \mathbb{E}\left[\max_{i=0,..,n} \mu^{n+1}(\theta^i) \middle| \mathcal{F}^n, \theta^n = \theta\right] - \max_{i=0,..,n} \mu^n(\theta^i)|_{\theta^n=\theta}. \tag{16}$$

In the Gaussian process regression framework, $\mu^{n+1}(\theta)$ given $\mathcal{F}^n$ is normally distributed for each value of θ, and the KGCP can be calculated exactly (see (Scott et al. 2011) or its description in Powell and Ryzhov (2012)[Chapter 16]). Thus, the next sampling decision is chosen to maximize the KGCP,

$$\theta^n \in \arg\max_\theta \bar{\nu}^{KG,n}(\theta).$$

After N observations, the implementation decision θ^n is chosen by maximizing the regression function,

$$\theta^\star \in \arg\max_\theta \mu^N(\theta).$$

The value of θ which maximizes equation (15) produces the best policy within the class of polices, $X^\pi(S_t|\theta)$. KGCP has been proven to converge asymptotically to the optimal value of θ (see Scott et al. (2011)).

6. Benchmark application: energy storage

Consider a power system as shown in Figure 2, involving an intermittent energy supply, an electricity demand, a simple interconnecting grid, and a battery storage device. At time t, the energy flows in Figure 2 are given by the vector

$$x_t \stackrel{\text{def}}{=} \left(x_t^{WR}, x_t^{GR}, x_t^{RD}, x_t^{WD}, x_t^{GD}\right).$$

Here, x_t^{IJ} denotes the amount of energy transferred from I to J at time step t. The superscript W stands for energy source (wind), D for demand, R for storage, and G for grid. All these entities are assumed to be nonnegative except x_t^{GR}. A negative value for x_t^{GR} refers to selling the electricity from the storage to the grid.

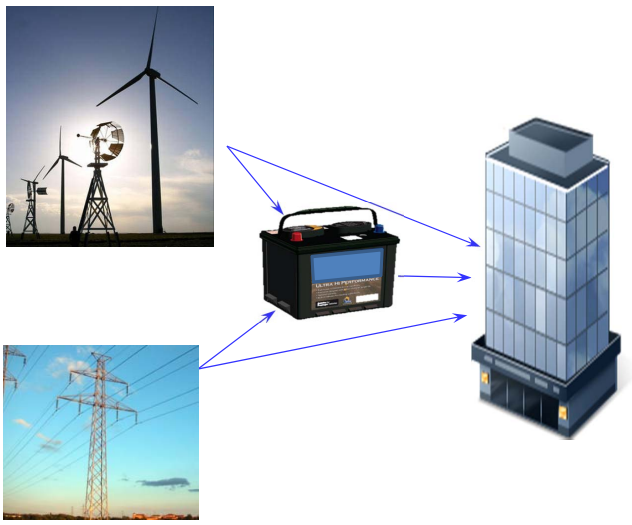

Figure 2 **The energy flow diagram.**

Denote the total electricity demand (in MWh) over the time period starting at $t-\Delta t$ and ending at t, by D_t. At every time step, the demand D_t must be served through either the wind energy, available electricity from the storage device, or energy purchased from the grid,

$$x_t^{GD} + \eta^{discharge} x_t^{RD} + x_t^{WD} = D_t. \tag{17}$$

Here, $0 < \eta^{discharge} < 1$ denotes the *discharging efficiency rate*.

The wind energy generated during the time period $[t-\Delta t,\ t)$, denoted by E_t, first flows to the electricity demand, i.e.,

$$x_t^{WD} = \min\{E_t,\ D_t\}. \tag{18}$$

The surplus is then sent to the storage device to be stored for the future use, i.e.,

$$x_t^{WR} + x_t^{WD} = E_t. \tag{19}$$

This constraint assumes that the storage device can dissipate some part of the power received from the wind source upon receiving, if it is more than its capacity.

Let R^{cap} indicate the total capacity of the storage. Define the constants ΔR^{min} and ΔR^{max} as the minimum and maximum fractions of the storage device that can be charged over Δt. For example, for a lead acid battery with a $C/10$ maximum charge and discharge rates, and $\Delta t = 15\text{min}$, $\Delta R^{min} = -1/40$ and $\Delta R^{max} = 1/40$. To avoid charging or discharging the storage device faster than the permitted rates, x_t^{GR} must satisfy:

$$\frac{\Delta R^{min} R^{cap}}{\eta^{discharge}} \leq x_t^{GR} \leq \frac{\Delta R^{max} R^{cap}}{\eta^{charge}}, \tag{20}$$

where η^{charge} is the *charging efficiency rate*. When selling the electricity from the storage to the grid is prohibited, the lower bound in equation (20) must be replaced by 0, i.e., $x_t^{GR} \geq 0$. Similarly, to ensure that the storage device is not discharged faster than allowed when sending energy from the storage device to the demand, we include:

$$0 \leq x_t^{RD} \leq \Delta R^{max} R^{cap}. \tag{21}$$

By implementing the energy flow X_t^π, the next storage state can be computed as

$$R_{t+\Delta t}(x_t) = \min\left\{ R_t + \frac{\eta^{charge}(x_t^{GR} + x_t^{WR}) - x_t^{RD}}{R^{cap}}, 1 \right\}. \tag{22}$$

In addition, we must have $R_{t+\Delta t} \geq R^{\min}$, where $R^{\min}$ is the minimum fraction of the capacity of the storage device that must remain full. For example, stationary lead-acid batteries with tubular plates, which are one of the lowest cost technologies for energy storage, should not be discharged below 20% of their capacity, i.e., $R^{\min} = 0.20$, e.g., see (Brunet 2011).

Constraints (19), (20), (21), (22), in addition to standard measurability conditions, together define the set of admissible policies. Note that equations (18), (19), and (17) effectively reduce the size of our action space from 5 dimensions to 2 dimensions.

The contribution function at every time step t then is the dollar value of energy sold minus the amount bought from the grid,

$$C(S_t, X_t^\pi) = P_t D_t - P_t \left(x_t^{GR} + x_t^{GD} \right). \tag{23}$$

The goal is to find a policy over the ergodic infinite horizon planning horizon which maximizes the accumulated expected discounted future rewards,

$$\max_\pi \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t C\left(S_t, X_t^\pi\right) \right].$$

In the absence of the wind source and the electricity load, in which the storage device is used solely to buy/sell the electricity from/to the grid, the model is referred to as the battery arbitrage problem.

Before seeking an optimal policy, it is necessary to define E_t, D_t, and prices P_t over time steps. Uncertainties in the energy supply (wind), electricity demand, and prices are modeled through stochastic processes, as described in the subsequent subsections.

6.1. Wind Power

The power output from the wind turbine over $[t,\ t+\Delta t)$, measured in MWh, is computed as below (see (MacKay 2009) page 263):

$$E_t = \frac{10^{-8}}{36} \times 0.5\ C_p\ \rho\ 50^2\pi\ W_t^3\ \Delta t. \tag{24}$$

Here, ρ is the density of air which is around $\rho = 1.225\ kg/m^3$. The term $50^2\pi\ m^2$ approximates the area in square meters swept by the rotor blades of the turbine. The power coefficient is around $C_p = 0.45$. In equation (24), W_t denotes the wind speed measured in meters per second and Δt is stated in seconds. Velocity of the wind, W_t, in m/s, equals

$$W_t = \left(Y_t + \mathbb{E}\left[\sqrt{W_t}\right]\right)^2.$$

where Y_t is the de-meaned square root of the wind speeds. Similar to Brown et al. (1984), we let Y_t evolve by an AR(1) model:

$$Y_t = \phi_1 Y_{t-\Delta t} + \epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon).$$

Using 15-min data from the wind speeds at Maryneal, Texas and applying the Yule-Walker equations, e.g., see (Carmona 2004) to fit the above model, we obtained $\mathbb{E}[\sqrt{W_t}] = 1.4781$, $\phi_1 = 0.7633$, $\sigma_\epsilon = 0.4020$. Note that other models have also been used in the literature; for example, Chen et al. (2010) suggest using a more general ARIMA model for the square root of the wind speed.

6.2. Electricity Prices

Figure 3 illustrates the electricity spot prices at the PJM Western Hub for one week. While the average of these real-time prices over 2009-2010 is \$42.11 per MWh, the prices are occasionally negative and have a maximum of \$362.90 per MWh. Figure 3 shows that the prices are lowest at night; they begin to increase around 5 am and are typically the highest in the evening around 6 pm.

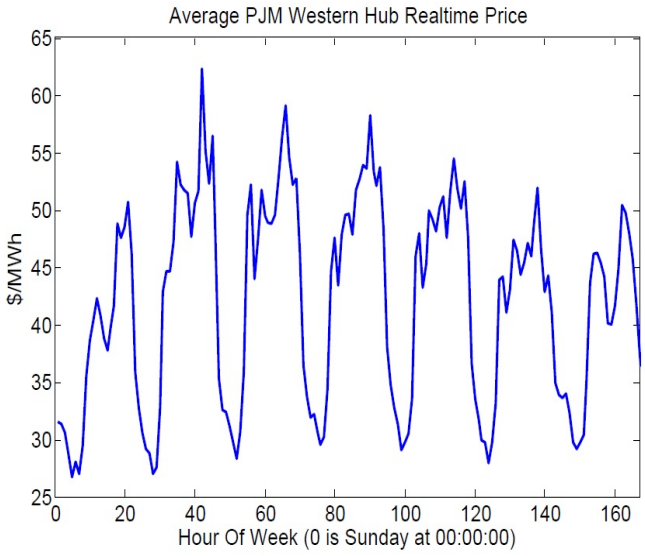

Figure 3 The average PJM Real-Time price at the Western Hub as a function of the hour of the week.

In addition to seasonality, mean reversion and heavy-tailed behaviour have been considered as important properties of energy prices, (see Eydeland and Wolyniec (2003)). Therefore, as in Cartea and Figueroa (2005), here we adopt a mean reverting jump diffusion model with seasonality for real-time electricity prices in $\$/MWh$,

$$P_t = \exp\left(Y_t^{ds} + Y_t^s\right) - c, \tag{25}$$

where the constant parameter c equals one minus the minimum value of P_t in the data set, i.e., $c = 1 - \min_t\ P_t$. In equation (25), the seasonal component Y_t^s is a deterministic periodic function of t, including an hour-of-week and a month-of-year components. The deseasonalized log prices are modeled using the following discretized mean reverting jump diffusion model:

$$Y_t^{ds} - Y_{t-\Delta t}^{ds} = \lambda(\mu - Y_{t-\Delta t}^{ds})\Delta t + \sigma\sqrt{\Delta t}\ \epsilon_t + J_t, \tag{26}$$

where μ is the long term equilibrium price, λ is the mean reversion rate, $\{\epsilon_{t+n\Delta t}\}_{n=0}^N$ are i.i.d. standard normal random variables, and J_t is the jump over the interval $[t-\Delta t,\ t)$. As in Cartea and Figueroa (2005), we model the jumps by the i.i.d. process,

$$J_t = \epsilon_t^{jump}\mathbf{1}\left(U_t < p^{jump}\right),$$

where $\epsilon_t^{jump} \sim N(0, \sigma^{jump})$ is the size of a jump, $U_t \sim \text{unif}(0,1)$, and p^{jump} is the probability of a jump over a time interval of length Δt. In addition, jump sizes $\{\epsilon_{t+n\Delta t}^{jump}\}_{n=0}^N$ are i.i.d. normal random variables.

We identify the nonzero jumps (as in Cartea and Figueroa (2005)) by counting the number of times where the absolute value of the return is more than three times its standard deviation. The probability of a jump p^{jump} equals then half of the fraction of times that jumps occur; we divide this estimate by two because most jumps are immediately followed by jumps in the opposite direction. The standard deviation of the magnitude of these jumps is considered as σ^{jump}. After estimating the jump terms, the parameters λ, μ, and σ can be estimated using least-squares linear regression on equation (26). For the electricity spot prices at the PJM Western Hub data set, we get $\sigma^{jump} = 0.4229$, $p^{jump} = 0.0170$, $\lambda = 1800.9$, $\mu = 4.1995$, $\sigma = 11.0971$, and $c = 27.2531$.

6.3. Electricity Demand

Eydeland and Wolyniec (2003) outline typical models for residential, commercial, and industrial power demands. While industrial power demand is relatively stable, residential power demand is highly dependent upon the temperature and exhibits seasonal variations. These hourly and daily features can be clearly observed from the left plot in Figure 4; the hourly ERCOT energy load starts ramping up in the morning around 5 am and peaks in the evening around 6 pm, although the patterns vary greatly based on the day of the week and the month of the year. Many approaches to forecast loads, such as an end-use model that incorporates appliances and customers, various regression models (based on temperature, time, and other factors), and heuristics made by experts, are summarized in Feinberg and Genethliou (2010). Similar to Pirrong and Jermakyan (2008) (and for the sake of simplicity) we adopt a Brownian motion with seasonality components for modeling the electricity demands in MWh,

$$D_t = m_t^{hour} + m_t^{month} + D_t^{ds}. \tag{27}$$

Here, m_t^{hour} and m_t^{month} indicate the hour-of-week seasonal and the month-of-year seasonal components, and D_t^{ds} is the deseasonalized load. The hour-of-week seasonal component m_t^{hour} is calculated as the average load over each of the hours of the week, and m_t^{month} is computed as the average load over each of the months of a year. The deseasonalized load D_t^{ds} evolves with temperature (in degrees Celsius) T_t at time t, as

$$D_t^{ds} = \alpha_0 + \alpha_1 T_t + \alpha_2 T_t^2 + \alpha_3 T_t^3 + \alpha_4 T_t^4 + \alpha_5 T_t^5 + \epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon) \tag{28}$$

where the coefficients $\alpha_0, \alpha_1, \cdots, \alpha_5$ can be estimated through least squares polynomial regression. The right plot in Figure 4 illustrates the estimated deseasonalized load D_t^{ds} against temperature. For the hourly ERCOT energy load data we obtained, $\phi_1 = 0.9636$, $\sigma_\epsilon^2 = 914870$.

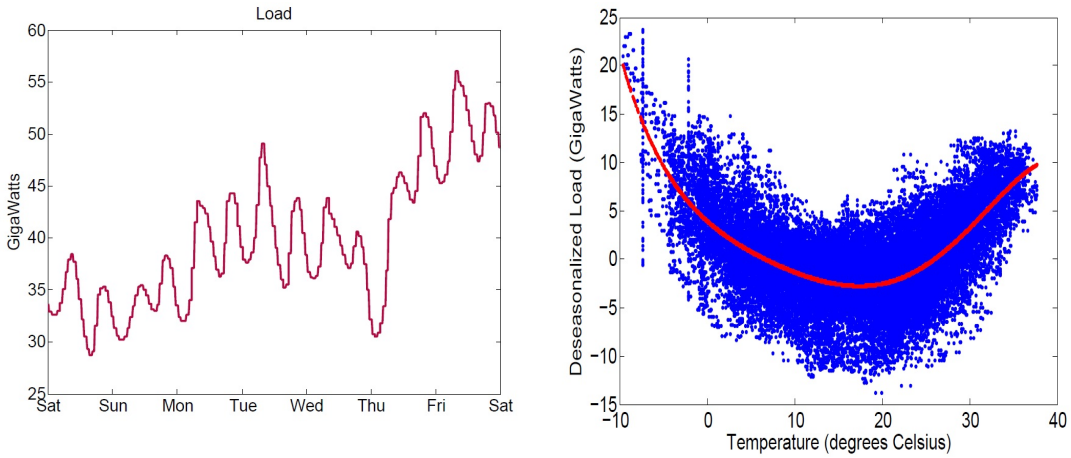

Figure 4 **Left plot: The total ERCOT load over the first full week of 2010. Right plot: Deseasonalized total ERCOT load versus temperature at Maryneal, TX over 2010. There is one point for each 15 minute period of 2010. The fitted fifth order polynomial regression is shown as well.**

6.4. Stochastic Dynamic Programming Formulation

We formulate this multi-stage stochastic programming problem through stochastic dynamic programming. The state variable, S_t, and the post-decision state variable, S_t^x, are defined as,

$$S_t = (R_t,\ E_t,\ D_t,\ P_t)\,,$$
$$S_t^x = (R_{t+\Delta t}(x),\ E_t,\ D_t,\ P_t)\,.$$

Recall that R_t is the fraction of the storage device that is full, E_t is the current amount of wind energy, D_t indicates the current energy demand, and P_t is the current spot price of electricity selling to (or purchasing from) the electrical grid. The *exogenous information process* is defined as the random changes in the state of the system, $W_{t+\Delta t} = \{\hat{E}_{t+\Delta t}, \hat{D}_{t+\Delta t}, \hat{P}_{t+\Delta t}\}$, which refer to exogenous changes in E_t, D_t and P_t. These exogenous changes may be state dependent as well as time dependent.

The goal is then to find the policy, $X^\pi(S_t)$, which maximizes the expected discounted future rewards,

$$\max_{\pi\in\Pi} \mathbb{E}\left[\sum_{t=0}^{\infty}\gamma^t\ C\left(S_t, X^\pi(S_t)\right)\right]. \tag{29}$$

The policy $X^\pi(S_t)$ may be stationary or time-dependent, which may be denoted by $X_t^\pi(S_t)$. Time-dependent policies can be explained by modifying the state variable through including time.

In the following section, we use the LSAPI algorithm, described in Section 2 with policy evaluation methods discussed in Section 3, as well as the direct policy search described in Section 5, for finding policies, which are then compared to optimal policies so we can evaluate their performance.

7. Numerical Experiments

Our main objective in this section is to assess the performance of the API variants and direct policy search. In subsection 7.1, these algorithms are compared with 20 discretized benchmark problems for energy storage operation management. The size of the benchmark problems considered are realistic relative to real world instances. Exact optimal stochastic dynamic programming solutions of these benchmark problems have been computed using the exact value iteration method. Finally, we run approximate policy iteration on a problem with a state consisting of five continuous dimensions and actions consisting of five continuous dimensions.

Throughout, the time step is fixed to $\Delta t = 15$ min. We use a discount factor, $\gamma = 99.90\%$. We found that discount factors of $\gamma = 99\%$ or smaller produce policies that are relatively myopic, and do not allow us to hold energy in storage for extended periods.

7.1. Benchmark Problems with Discrete State Spaces

We first consider finite, discretized state and action spaces with a fixed probability transition matrix. An exact solution for an infinite horizon problem can be found using the value iteration method, e.g., see Puterman (1994). We note that computing these optimal policies typically required approximately two *weeks* of CPU time, due primarily to the very high discount factor. In this method, $V^0(s)$ is initialized to a constant for all states s in

the state space, and at each iteration n, the algorithm updates the value function at each state using

$$V^n(s) = \max_x \left\{ C(s,x) + \gamma \sum_{s' \in \mathcal{S}} V^{n-1}(s') \Pr(s'|s,x) \right\}, \quad \text{for every } s \in \mathcal{S}.$$

The algorithm will converge to within ϵ of the true value functions (see Puterman (1994)).

To establish a set of benchmark problems, we considered 20 instances of the energy storage operation management problem in Section 6. These problems are summarized in Table 1, where problem type "Full" refers to the problem in Figure 2 with energy from wind source and the grid serves an electricity demand. The problem type "BA" refers to a battery arbitrage problem, which considers only trading between the storage and the grid to take advantage of price variations. We discretized the state space in the benchmark test problems and then created fixed probability transition matrices for the exogenous information process in order to create a true discrete process. Table 1 also reports how finely each state variable is discretized (the size of the state space for a particular problem is the product of each of the discretization levels). We then list the average maximum wind capacity divided by the load, the storage capacity divided by the load over an hour, the round trip efficiency (RTE) of the storage device, and the max charge and discharge rate of the storage device. For example, $C/10$ indicates that the storage device can be fully charged or discharged within 10 hours. The transition matrix of the electricity prices was fitted using the PJM Western Hub real time prices (with and without time of day). The transition matrix of the load was fitted using the load of the PJM Mid-Atlantic Region (with time of day). The transition matrix for the wind was fit using data from wind speeds near the Sweetwater Wind Farm. For Problems $1-16$ the state space is resource level, wind energy, and electricity price, i.e., $S_t = (R_t, E_t, P_t)$. For this experiment, time and demand are held fixed in order to keep the benchmark problems computationally tractable, as the exact value iteration, even for these simplified problems, requires approximately two weeks on a 3Ghz processor. For Problems $17-20$, the state space is the time-of-day, τ_t, (1-96 corresponding to fifteen minute intervals in a day), the resource level, and the electricity price. The full state variable is then given by $S_t = (\tau_t, R_t, P_t)$.

To specify reasonable values for the maximum number of policy evaluation and policy improvement iterations, M and N in Figure 1, we implemented the LSAPI method using

Problem		Number of Discretization Levels					Parameters			
Number	Type	Time τ_t	Resource R_t	Price P_t	Demand D_t	Wind E_t	Wind $\frac{E_t}{D_t}$	Storage $\frac{R_t}{D_t}$	RTE	Charge Rate
1	Full	1	33	20	1	10	0.1	2.5	.81	C/10
2	Full	1	33	20	1	10	0.1	2.5	.81	C/1
3	Full	1	33	20	1	10	0.1	2.5	.70	C/10
4	Full	1	33	20	1	10	0.1	2.5	.70	C/1
5	Full	1	33	20	1	10	0.2	2.5	.81	C/10
6	Full	1	33	20	1	10	0.2	2.5	.81	C/1
7	Full	1	33	20	1	10	0.2	2.5	.70	C/10
8	Full	1	33	20	1	10	0.2	2.5	.70	C/1
9	Full	1	33	20	1	10	0.1	5.0	.81	C/10
10	Full	1	33	20	1	10	0.1	5.0	.81	C/1
11	Full	1	33	20	1	10	0.1	5.0	.70	C/10
12	Full	1	33	20	1	10	0.1	5.0	.70	C/1
13	Full	1	33	20	1	10	0.2	5.0	.81	C/10
14	Full	1	33	20	1	10	0.2	5.0	.81	C/1
15	Full	1	33	20	1	10	0.2	5.0	.70	C/10
16	Full	1	33	20	1	1	0.2	5.0	.70	C/1
17	BA	96	33	20	1	1	-	-	.81	C/10
18	BA	96	33	20	1	1	-	-	.81	C/1
19	BA	96	33	20	1	1	-	-	.70	C/10
20	BA	96	33	20	1	1	-	-	.70	C/1

Table 1 Benchmark problems: the number of discretization levels for time (1=steady state), resource, price, load (1=deterministic) and wind.

instrumental variables Bellman error minimization on the 17^{th} benchmark problem several times. For this test problem, as illustrated in Figure 5, most of the improvement has occurred before $N = 5000$ iterations of policy evaluations and $M = 30$ iterations of policy improvement step. Thus, in the rest of this section, we fix $N = 5000$ and $M = 30$.

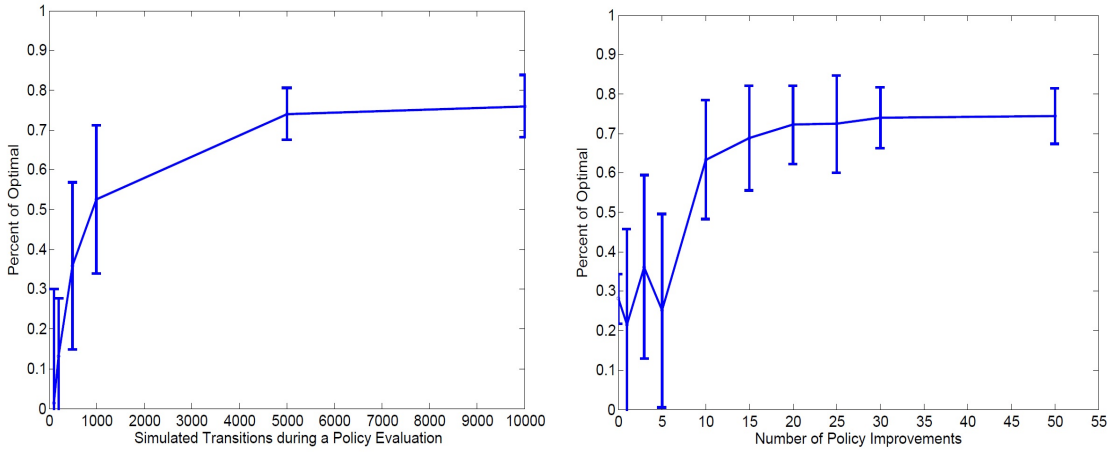
Percent of Optimal
Simulated Transitions during a Policy Evaluation
Percent of Optimal
Number of Policy Improvements

Figure 5 Left Plot: Progress of the LSAPI method using instrumental variable Bellman error minimization for benchmark problem 17 as N varies. Right Plot: Progress of the LSAPI method using instrumental variable Bellman error minimization for benchmark problem 17 as M varies.

7.2. API and Direct Policy Search Policies for the Benchmark Problems

Figure 6 compares the LSAPI method with instrumental variables Bellman error minimization (IVAPI), API with least-squares Bellman error minimization (LSAPI), the myopic policy (Myopic), and direct policy search based on KGCP (Direct). In this figure, quadratic basis functions are used to implement the IVAPI and LSAPI methods. The myopic policy discharges the storage device as quickly as possible and keeps it at its minimum level since then. The value of the myopic policy is still positive due to the wind source as well as the initial storage level.

We run each algorithm 100 times. For each run of the algorithms, the final policies produced by each algorithm are then evaluated on the same sample paths, $\omega \in \Omega$, where ω is generated from the discretized exogenous information process. We then record the average percent of optimal and the standard deviation of the average percent of optimal across the 100 runs. The average percentage of optimal for a policy π is computed as

$$\% \text{ of optimal} = \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \frac{\hat{F}^{\pi}\left(S_0(\omega)\right)}{V^{\star}\left(S_0(\omega)\right)},$$

where ω is a sample path of the randomness in the state transitions, and $S_0(\omega)$ is the starting state which has been randomly generated from a uniform distribution. $\hat{F}^{\pi}(S_0(\omega))$ is a realization of value of the policy π run on the sample path ω, starting at the state $S_0(\omega)$, and $V^{\star}(S_0(\omega))$ is the true value of the optimal policy for state $S_0(\omega)$ computed using the exact value iteration method.

Similarly, we implemented direct policy search using KGCP 100 times, and computed the average percent of optimal and its standard deviation. To implement direct policy search using KGCP, we budget ourselves to simulating 50 sequentially chosen policies, after which the KGCP algorithm must choose what it believes to be the best policy.

Figure 6 illustrates that IVAPI significantly outperforms LSAPI, but still underperforms the optimal policy by a wide margin. Direct policy search produced solutions that are on average 91.80 percent of optimal, and are always at least 70% of optimal for problems 1 through 16. This suggests that for the benchmark problems on the application of interest, direct policy search is much more robust relative to the LSAPI methods. However, direct policy search quickly becomes intractable as the number of basis functions increases. In addition, choosing the search domain for direct policy search is another significant complication as the number of basis functions increases. We suggest using LSAPI to find good

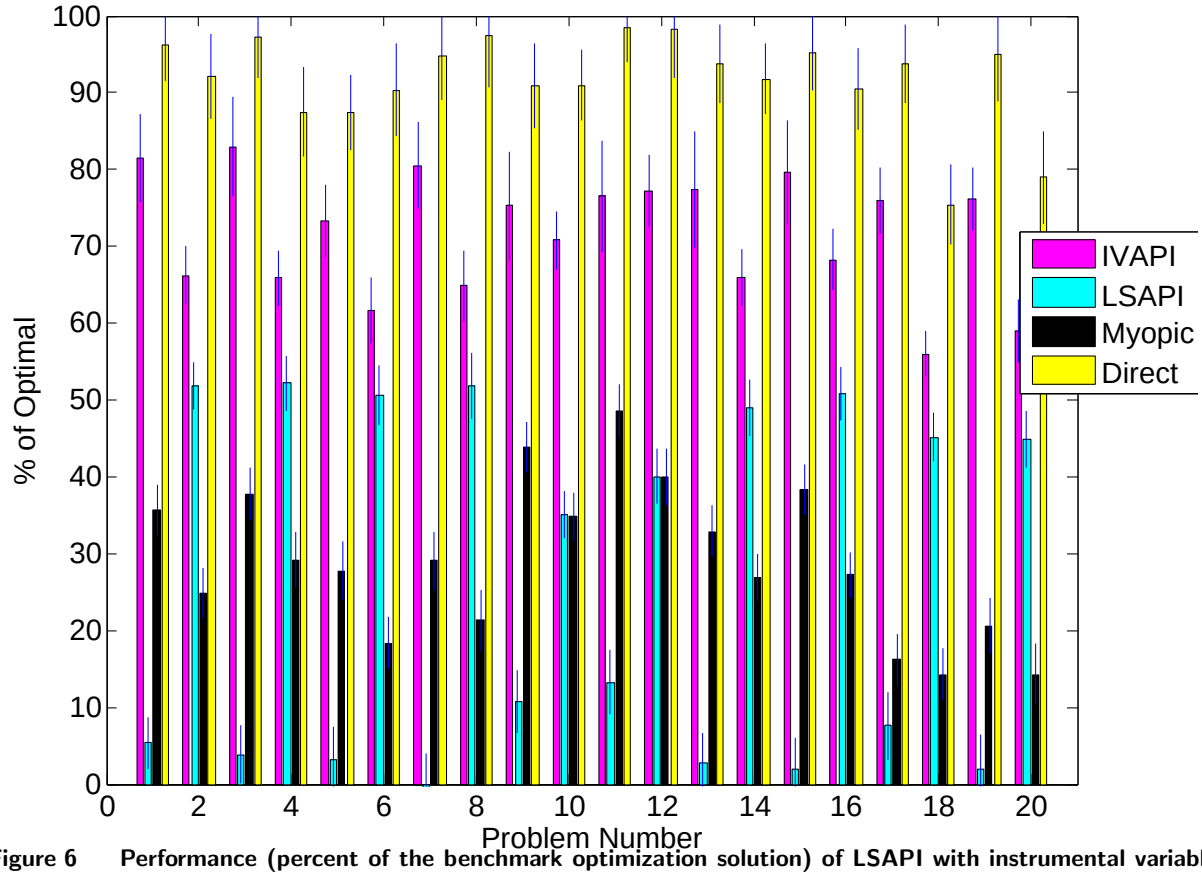

Figure 6 **Performance (percent of the benchmark optimization solution) of LSAPI with instrumental variables, LSAPI, the myopic policy and direct policy search.**

values of the regression parameters, and then apply direct policy search to improve the policy in the region of the fitted regression parameters.

Figure 7 depicts a sample path of an IVAPI policy for the 1^{st} benchmark problem in Table 1. The storage device is charged when electricity prices are low and discharged when electricity prices are high. We also note that the storage device fully discharges (below 20%) relatively infrequently.

In order to reduce the number of basis functions in the algorithms, one may consider smaller dimensions for the post-decision state when constructing the value function approximation. Figure 8 shows the results using three value function approximations: (1) all three state variables (R_t, E_t, and P_t), (2) only R_t, and (3) the resource and electricity prices R_t and P_t. It is observed that using the resource level alone for the domain of the post-decision value function results in quite poor performances for most benchmark problems. Using both R_t and P_t appears to do fairly well overall, although using all of the state variable dimensions yields the best results. For certain problems, it may actually be advantageous

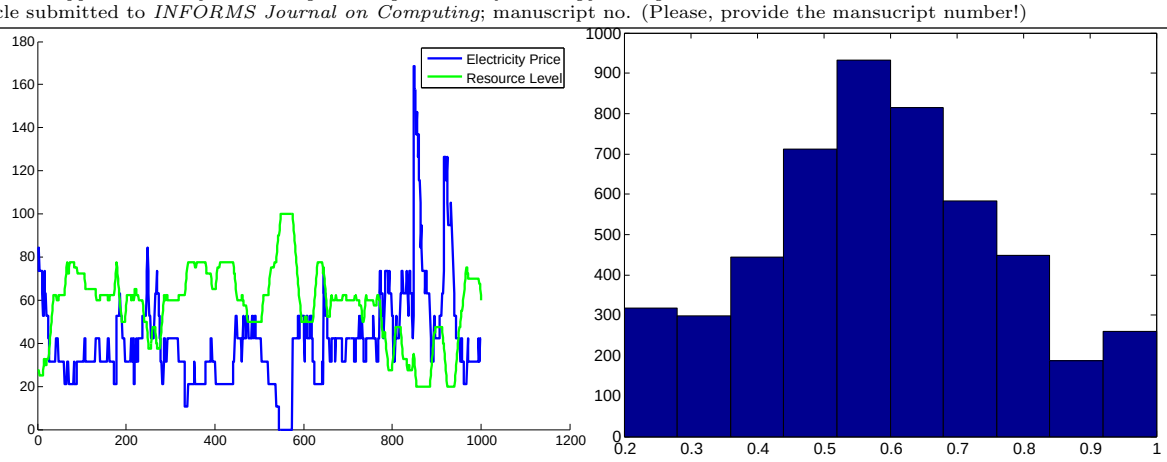

(a) (b)

Figure 7 **We plot a 10 day sample path of an IVAPI policy using quadratic basis functions on the 1^{st} benchmark problem. Plot (a): The electricity price and resource level R_t. Plot (b): Histogram of the resource level.**

to leave variables out of the value function approximation to simplify the policy search process.

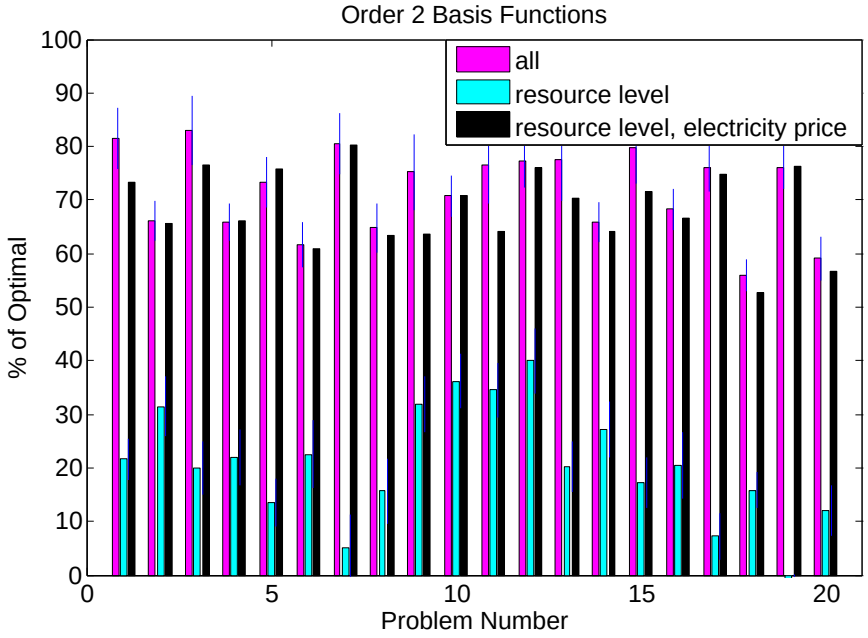

Figure 8 **The percentage of optimal along with 95% confidence intervals for the average percent of optimal for the IVAPI policy, derived using quadratic basis functions, when only certain dimensions of the post-state are included in the post-state value function approximation.**

7.3. Benchmark Problems with Continuous State Spaces

In this section we consider a set of 10 problems with continuous state spaces, continuous actions, and the state transitions corresponding to the problem discussed in Section 6. These problems are listed in Table 2. The electricity prices and loads are now time-dependent and stochastic for Problems 1-3. Problems 4-10 are continuous steady-state problems. We compare both API algorithms and the myopic policy, although an optimal solution will no longer be available.

Problem		Number of Discretization Levels					Parameters			
Number	Type	Time	Resource	Price	Load	Wind	Wind	Storage	RTE	Charge Rate
1	Full	96	Cont.	Cont.	Cont.	Cont.	0.1	2.5	.81	C/10
2	Full	96	Cont.	Cont.	Cont.	Cont.	0.1	5.0	.81	C/10
3	BA	96	Cont.	Cont.	1	1	-	-	.81	C/10
4	Full	1	Cont.	Cont.	Cont.	Cont.	0.1	5.0	.81	C/10
5	Full	1	Cont.	Cont.	Cont.	Cont.	0.1	2.5	.81	C/1
6	Full	1	Cont.	Cont.	Cont.	Cont.	0.1	2.5	.70	C/1
7	BA	1	Cont.	Cont.	1	1	-	-	.81	C/10
8	Full	1	Cont.	Cont.	Cont.	Cont.	0.1	5.0	.81	C/1
9	Full	1	Cont.	Cont.	Cont.	Cont.	0.1	5.0	.70	C/1
10	Full	1	Cont.	Cont.	Cont.	Cont.	0.2	2.5	.81	C/1

Table 2 Parameter settings for problems with continuous states. Problems 1-3 have time-dependent stochastic loads and electricity prices. Problems 4-10 are steady-state.

Figure 9(a) shows that for all 10 test problems described in Table 2, IVAPI prominently outperform LSAPI and the myopic policy. We see again that the use of instrumental variables in the API method adds value. Although all the dimensions of the state variable and action space are difficult to visualize, in Figure 9(b) we use an IVAPI policy to show the electricity price and R_t on one particular sample path. This figure shows that the policy tends to start charging the battery at night when electricity prices are low and then discharges the storage device throughout the day when electricity prices are higher.

The results in this section suggests that IVAPI is a promising and scalable algorithm, suitable for problems where the states are continuous and the transition probabilities are unknown.

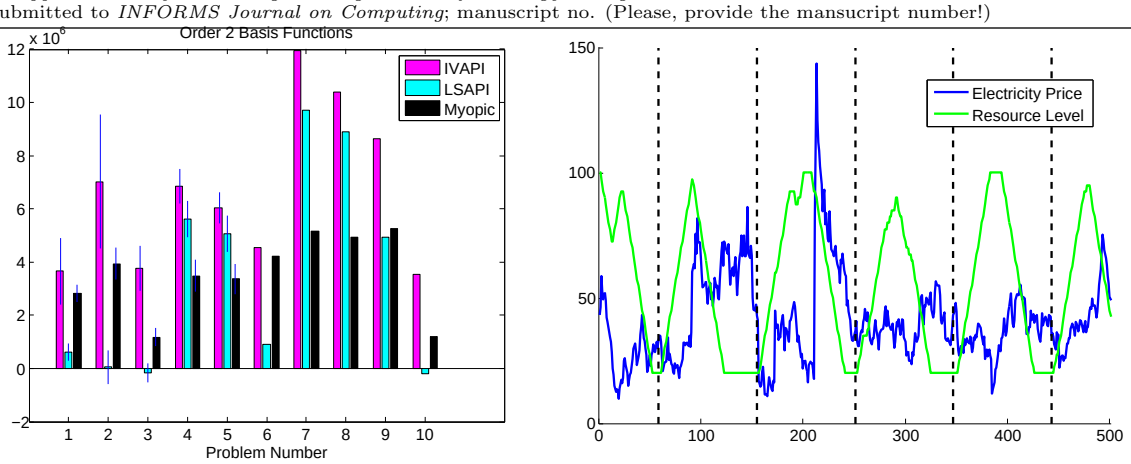

(a) (b)

Figure 9 **Plot (a): Average objective of the myopic, IVAPI, and LSAPI policies on the continuous problems described in Table 2. Plot (b): A sample path of the electricity spot price and resource level. The vertical lines correspond to midnight.**

8. Conclusions

In this paper, we have studied four variants of approximate policy iteration (API) methods, based on Bellman error minimization policy evaluation. We have considered least-squares Bellman error minimization, Bellman error minimization using instrumental variables, least-squares projected Bellman error minimization, and projected Bellman error minimization using instrumental variables. We have established the equivalence between Bellman error minimization using instrumental variables and projected Bellman error minimization methods. This result unifies algorithmic strategies that were previously thought to be fundamentally different.

The approximate policy iteration methods were then evaluated numerically. Motivated by control problems in energy storage applications, we created a series of benchmark problems to compare the different algorithmic strategies. The performance of a policy search method was also evaluated on those benchmark problems.

We can draw several interesting conclusions from our numerical work. Bellman error minimization using instrumental variables appears to improve significantly over the basic LSAPI method, but otherwise did not work well when compared to the optimal benchmark, producing results that ranged between 60 and 80 percent of optimal. Direct policy search performed much better, with results averaging over 90 percent of optimal. Given that this is a problem that is ideally suited to LSAPI, it calls into question whether this is a method that can be counted on to produce good results. The difficulty can probably be traced

to the need to use pure exploration when sampling the state space, but given that the value function is relatively smooth and well approximated using a quadratic, it seems a bit surprising that the use of pure exploration would produce errors that are this large.

This research suggests that there are clear advantages to using direct policy search, possibly in conjunction with approximate policy iteration. The challenge is that in its derivative-free form, policy search does not scale easily with the dimension of the parameter space. This may be a major limitation in time-dependent applications, where we may need to estimate a different set of parameters for each time period.

Appendix A: The Instrumental Variable Method

The instrumental variable method (IVM) is a well known technique for dealing with errors in the explanatory variables of a regression problem, and provides a way to obtain consistent parameter estimates, see e.g., (Durbin 1954, Bradtke and Barto 1996, Young 2011, Kendall and Stuart 1961, Cameron and Trivedi 2005). Consider the linear model in the matrix form:

$$Y = X\beta, \tag{30}$$

where Y is a $n \times 1$ vector of response variables, X is a $n \times k$ matrix of explanatory variables, and β is a $k \times 1$ vector of weights. Let X' and Y' be observable values of the true values X and Y. Following similar notation as in (Durbin 1954), denote the errors in the observed values of X and Y by X'' and Y'', respectively.

$$X' = X + X'', \tag{31}$$

$$Y' = Y + Y''. \tag{32}$$

Hence, the linear regression model can now be written as $Y' = X\beta + Y''$. The least-squares method is probably the most widely used technique to estimate β. Kendall and Stuart (1961) and Durbin (1954) show that least squares regression encounters problems with the model given by equations (30), (31), (32), since the observation of X is typically correlated with the error in X. If β is a scalar ($k = 1$), this is easy to show. Starting with the least squares estimate of β_1,

$$\begin{aligned}
\hat{\beta}_1 = ((X')^\top X')^{-1}(X')^\top Y' &= \frac{\sum_{i=1}^n X'_{i1} Y'_i}{\sum_{i=1}^n (X'_{i1})^2} \\
&= \frac{\sum_{i=1}^n X'_{i1}(X_{i1}\beta_1 + Y''_i)}{\sum_{i=1}^n (X'_{i1})^2} && (33) \\
&= \frac{\sum_{i=1}^n X'_{i1}\left((X'_{i1} - X''_{i1})\beta_1 + Y''_i\right)}{\sum_{i=1}^n (X'_{i1})^2} && (34) \\
&= \beta_1 - \beta_1 \frac{\sum_{i=1}^n X'_{i1} X''_{i1}}{\sum_{i=1}^n (X'_{i1})^2} + \frac{\sum_{i=1}^n X'_{i1} Y''_i}{\sum_{i=1}^n (X'_{i1})^2}, && (35)
\end{aligned}$$

where equations (30) and (32) were used in equation (33), and equation (34) follows (31). By taking the limit as $n \to \infty$ and assuming $\mathrm{Cov}[X'_{i1}, Y''_i] = 0$, equation (35) becomes

$$\lim_{n\to\infty} \hat{\beta}_1 = \beta_1 - \beta_1 \lim_{n\to\infty} \left(\frac{\sum_{i=1}^n X'_{i1} X''_{i1}}{\sum_{i=1}^n (X'_{i1})^2} \right) + \lim_{n\to\infty} \left(\frac{\sum_{i=1}^n X'_{i1} Y''_i}{\sum_{i=1}^n (X'_{i1})^2} \right)$$

$$= \beta_1 - \beta_1 \lim_{n\to\infty} \left(\frac{\sum_{i=1}^n X'_{i1} X''_{i1}}{\sum_{i=1}^n (X'_{i1})^2} \right) + \frac{\mathrm{Cov}[X'_{i1}, Y''_i]}{\mathbb{E}[(X'_{i1})^2]}$$
$$= \beta_1 - \beta_1 \lim_{n\to\infty} \left(\frac{\sum_{i=1}^n X'_{i1} X''_{i1}}{\sum_{i=1}^n (X'_{i1})^2} \right). \tag{36}$$

For many problems, X'_{i1} and X''_{i1} are positively correlated. Whence equation (36) implies that the least squares estimate $\hat{\beta}_1$ for β_1 is inconsistent. However, a consistent estimate can be obtained when an instrumental variable is available.

As it is discussed in the following section, a properly chosen instrumental variable can yield a consistent estimator for β. The following assumptions are made on the noise in X and Y:

Assumption A.1 $\mathbb{E}[Y''_i] = 0$, *for every* $i = 1, \cdots, n$.

Assumption A.2 $\mathbb{E}[X''_{ij}] = 0$, *for every* $i = 1, \cdots, n$, *and* $j = 1, \cdots, k$.

Unlike a standard linear regression problem in which a functional relationship to relate mean of the dependant variable to the regressor variable is considered, the IV method studies equation (30) as a structural relation between the two random variables X and Y, e.g., see Durbin (1954), Kendall and Stuart (1961).

Assume that an instrumental variable, Z_j, exists such that it is correlated with the true X_j but uncorrelated with the errors in the observations of X and Y.

Assumption A.3 $Cov[Z_{ij}, Y''_i] = 0$, *for every* $i = 1, \cdots, n$, *and* $j = 1, \cdots, k$.

Assumption A.4 $Cov[Z_{ij}, X''_{il}] = 0$, *for every* $i = 1, \cdots, n$, *and* $j, l = 1, \cdots, k$.

Assumption A.5 *The* $k \times k$ *matrix* Σ *has full rank* k, *where* X_l *indicate the* l*'th column of* X *and* $\Sigma_{jl} = Cov[Z_j, X_l]$.

Assumption A.6 $\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^n Z_{ij} Y''_i = 0$, *for every* $j = 1, \cdots, k$.

Assumption A.7 $\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^n Z_{ij} X''_{il} = 0$, *for every* $j, l = 1, \cdots, k$.

Assumption A.8 $\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^n Z_{ij} X_{il} = Cov[Z_j, X_l]$, *for every* $j, l = 1, \cdots, k$.

Assumptions A.6, A.7, and A.8 do not follow trivially from assumptions A.3, A.4, A.5 without additional assumptions such as independence across the n observations (because the law of large numbers does not apply). The method of instrumental variables defines the estimator $\hat{\beta}$ as the solution to

$$Z^\top X' \hat{\beta} = Z^\top Y', \tag{37}$$

where Z is a $n \times k$ matrix. Note that $\hat{\beta}$ is uniquely defined when $Z^\top X'$ has full rank k.

Below we extend the consistency proof from Kendall and Stuart (1961) to use multiple instrumental variables ($k > 1$).

Proposition A.1 *For the model given by equations* (30), (31), (32), *if assumptions A.1, A.2, A.3, A.4, A.5, A.6, A.7, A.8 hold, then* $\hat{\beta} = [Z^\top X']^{-1} Z^\top Y'$ *is a consistent estimator of* β.

Proof: *Using equation* (37) *and equations* (30), (31), *and* (32), *we get*

$$(Z^\top X + Z^\top X'')\hat{\beta} = Z^\top (X + X'')\hat{\beta} = Z^\top (X\beta + Y'') = Z^\top X\beta + Z^\top Y''. \tag{38}$$

Taking the limit as n goes to infinity yields,

$$\lim_{n\to\infty} \frac{1}{n} \left(Z^\top X + Z^\top X'' \right) \hat{\beta} = \lim_{n\to\infty} \frac{1}{n} \left(Z^\top X\beta + Z^\top Y'' \right).$$

By applying assumptions A.3 and A.6 which imply $\lim_{n\to\infty} \frac{1}{n} Z^\top Y'' = \vec{0}$ *and assumptions A.4 and A.7 which imply* $\lim_{n\to\infty} \frac{1}{n} Z^\top X'' = \mathbf{0}$, *and Slutsky's theorem when taking the limit, we arrive at:*

$$\lim_{n\to\infty} \frac{1}{n} Z^\top X\hat{\beta} = \lim_{n\to\infty} \frac{1}{n} Z^\top X\beta.$$

Since the sample covariances converge in probability to the true covariances by Assumption A.8, we get $\Sigma \left(\lim_{n\to\infty} \hat{\beta} \right) = \Sigma\beta$. *Now, Assumption A.5 which ensures that* $\hat{\beta}$ *is unique yields* $\lim_{n\to\infty} \hat{\beta} = \beta$. *Q.E.D.*

References

Barton, John P., David G. Infield. 2004. Energy storage and its use with intermittent renewable energy. *IEEE Transactions On Energy Conversion* **19**(2) 441–448.

Bertsekas, D. P., John Tsitsiklis. 1996a. *Neuro-Dynamic Programming*. 1st ed. Athena Scientific.

Bertsekas, Dimitri P. 2012. *Dynamic Programming and Optimal Control, Vol. II, 4th Edition: Approximate Dynamic Programming*. Athena Scientific, Belmont, MA.

Bertsekas, Dimitri P., John N. Tsitsiklis. 1996b. *Neuro-dynamic programming*. Athena Scientific, Belmont, MA.

Bertsekas, D.P. 2011. *Dynamic Programming and Optimal Control 3rd Edition, Volume II*. CRC.

Bowden, R.J., D.A. Turkington. 1984. *Instrumental variables*. Cambridge University Press.

Bradtke, Steven J., Andrew G. Barto. 1996. Linear least-squares algorithms for temporal difference learning. *Machine Learning* **22** 33–57.

Brown, Barbara G., Richard W. Katz, Allan H. Murphy. 1984. Time series models to simulate and forcast wind speed and power. *Journal of Climate and Applied Meteorology* **23** 1184–1195.

Brunet, Yves. 2011. *Energy Storage*. Wiley.

Cameron, A. Colin, Pravin K. Trivedi. 2005. *Microeconometrics: Methods and Applications*. Cambridge University Press.

Carmona, R. 2004. *Statistical analysis of financial data in S-Plus*. Springer Verlag.

Carmona, R., M. Ludkovski. 2005. Gas storage and supply guarantees: an optimal switching approach. *submitted to Management Science* .

Cartea, A., M.G. Figueroa. 2005. Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* **12**(4) 313–335.

Chen, Peiyuan, T. Pedersen, B. Bak-Jensen, Zhe Chen. 2010. Arima-based time series model of stochastic wind power generation. *IEEE Transactions on Power Systems* **25**(2) 667 –676.

de Farias, D. P., B. Van Roy. 2000. On the existence of fixed points for approximate value iteration and temporal-difference learning. *Journal of Optimization Theory and Applications* **105** 1201–1225.

De Farias, D.P., B. Van Roy. 2000. On the existence of fixed points for approximate value iteration and temporal-difference learning. *Journal of Optimization theory and Applications* **105**(3) 589–608.

Durbin, J. 1954. Errors in variables. *Revue de l'Institut International de Statistique* 23–32.

Eydeland, A., K. Wolyniec. 2003. *Energy and power risk management: New developments in modeling, pricing, and hedging*. John Wiley & Sons Inc.

Feinberg, E.A., D. Genethliou. 2010. Load forecasting. Joe H. Chow, Felix F. Wu, James A. Momoh, eds., *Applied Mathematics for Restructured Electric Power Systems: Optimization, Control, and Computational Intelligence*. Springer, 269–285.

Gonzalez, Javier Garcia, Rocio Moraga Ruiz de la Muela, Luz Matres Santos, Alicia Mateo Gonzalez. 2008. Stochastic joint optimization of wind generation and pumped-storage units in an electricity market. *IEEE Transactions on Power Systems* **23**(2) 460–468.

Greenblatt, J.B., S. Succar, D.C. Denkenberger, R.H. Williams, R.H. Socolow. 2007. Baseload wind energy: modeling the competition between gas turbines and compressed air energy storage for supplemental generation. *Energy Policy* **35**(3) 1474–1492.

Harsha, Pavithra, Munther Dahleh. 2012. Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy. *submitted* 1–34.

Kendall, M.G., A. Stuart. 1961. *The Advanced Theory of Statistics: Inference and Relationship*, vol. 2. Hafner Publishing Company.

Lagoudakis, M.G., R. Parr. 2003a. Least-squares policy iteration. *Journal of Machine Learning Research* **4** 1107–1149. URL `http://portal.acm.org/citation.cfm?id=945365.964290`.

Lagoudakis, M.G., R. Parr. 2003b. Least-squares policy iteration. *Journal of Machine Learning Research* **4** 1107–1149.

Lai, G., F. Margot, N. Secomandi. 2010. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations research* **58**(3) 564–582.

Ma, Jun, Warren B. Powell. 2010. Convergence analysis of kernel-based on-policy approximate policy iteration algorithms for Markov decision processes with continuous, multidimensional states and actions. *Working Paper* .

MacKay, David JC. 2009. *Sustainable Energy-Without the Hot Air*. 1st ed. UIT Cambridge Ltd.

Maei, Hamid H.R., Csaba Szepesvari, Shalabh Bhatnagar, David Silver, Doina Precup, Richard S. Sutton. 2009. Convergent Temporal-Difference Learning with Arbitrary Smooth Function Approximation. *Advances in Neural Information Processing Systems 22*. MIT Press, Vancouver, BC, 1–9. URL `http://books.nips.cc/papers/files/nips22/NIPS2009\_1121\_slide.pdfhttp://scholar.google.com/scholar?hl=en\&btnG=Search\&q=intitle:Convergent+Temporal-Difference+Learning+with+Arbitrary+Smooth+Function+Approximation\#0`.

Maei, Hamid H.R., Csaba Szepesvari, Shalabh Bhatnagar, Richard S. Sutton. 2010. Toward Off-Policy Learning Control with Function Approximation. *ICML-2010*. MIT Press.

Nascimento, J. M., Warren B. Powell. 2013. An Optimal Approximate Dynamic Programming Algorithm for Concave , Scalar Storage Problems with Vector-Valued Controls. *IEEE Transactions on Automatic Control* 1–16.

Pirrong, Craig, Martin Jermakyan. 2008. The price of power: The valuation of power and weather derivatives. *Journal of Banking & Finance* **32** 2520–2529.

Powell, Warren B. 2011. *Approximate Dynamic Programming: Solving the curses of dimensionality*. 2nd ed. John Wiley and Sons, New York.

Powell, Warren B., A. George, A Lamont, J Stewart. 2011. SMART: A Stochastic Multiscale Model for the Analysis of Energy Resources, Technology and Policy. *Informs J. on Computing* .

Powell, Warren B., Jun Ma. 2011. A review of stochastic algorithms with continuous value function approximation and some new approximate policy iteration algorithms for multidimensional continuous applications. *J. Control Theory and Applications* **9**(3) 336–352. doi:10.1007/s11768-011-0313-y.

Powell, Warren B., Ilya O. Ryzhov. 2012. *Optimal Learning*. John Wiley & Sons, Hoboken, NJ.

Puterman, M. L. 1994. *Markov Decision Processes*. John Wiley & Sons, New York.

Scott, Warren R., Peter Frazier, Warren B. Powell. 2011. The correlated knowledge gradient for simulation optimization of continuous parameters using Gaussian process regression. *SIAM Journal on Optimization* **21**(3) 996–1026.

Simao, Hugo P., J Day, A. George, T. Gifford, Warren B. Powell, J Nienow. 2009. An Approximate Dynamic Programming Algorithm for Large-Scale Fleet Management: A Case Application. *Transportation Science* **43**(2) 178–197.

Sioshansi, R., P. Denholm, T. Jenkin, J. Weiss. 2009. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics* **31**(2) 269–277.

Söderström, T., P. Stoica. 1983. *Instrumental variable methods for system identification*, vol. 161. Springer-Verlag Berlin.

Sørensen, B. 1981. A combined wind and hydro power system. *Energy Policy* **9**(1) 51–55.

Sutton, Richard S., Hamid H.R. Maei, Doina Precup, Shalabh Bhatnagar, David Silver, Csaba Szepesvari, Eric Wiewiora. 2009a. Fast gradient-descent methods for temporal-difference learning with linear function approximation. *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*. ACM Press, New York, New York, USA, 1–8. doi:10.1145/1553374.1553501. URL `http://portal.acm.org/citation.cfm?doid=1553374.1553501`.

Sutton, Richard S., Csaba Szepesvari, Hamid H.R. Maei. 2009b. A Convergent O (n) Algorithm for Off-policy Temporal-difference Learning with Linear Function Approximation. *Advances in Neural Information Processing Systems*, vol. 21. Citeseer, 1609–1616. URL `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.144.6298\&rep=rep1\&type=pdf`.

Sutton, R.S., A.G. Barto. 1998. *Reinforcement Learning*. The MIT Press, Cambridge, Massachusetts.

Sutton, R.S., H.R. Maei, D. Precup, S. Bhatnagar, D. Silver, C. Szepesvári, E. Wiewiora. 2009c. Fast gradient-descent methods for temporal-difference learning with linear function approximation. *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 993–1000.

Swider, D.J. 2007. Compressed air energy storage in an electricity system with significant wind power generation. *IEEE transactions on energy conversion* **22**(1) 95.

Szepesvari, Csaba. 2010. *Algorithms for Reinforcement Learning*, vol. 4. Morgan and Claypool. doi:10.2200/S00268ED1V01Y201005AIM009. URL `http://www.morganclaypool.com/doi/abs/10.2200/S00268ED1V01Y201005AIM009`.

Topaloglu, H., W. B. Powell. 2006. Dynamic programming approximations for stochastic, time-staged integer multicommodity flow problems. *Informs Journal on Computing* **18**(1) 31–42.

Tsitsiklis, J., B. Van Roy. 1997. An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control* **42**(5) 674–690.

Tsitsiklis, John N. 1994. Asynchronous stochastic approximation and Q-learning. *Machine Learning* **16** 185–202.

Xi, Xiaomin, Ramteen Sioshansi, Vincenzo Marano. 2014. A Stochastic Dynamic Programming Model for Co-optimization of Distributed Energy. *Energy Sytstems* .

Young, Peter C. 2011. *Recursive Estimation and Time-Series Analysis, 2nd ed*. Springer-Verlag, Berlin, Heidelberg.